

AMENDMENTS TO THE CLAIMS:

1. (Currently amended) A method of improving at least one of speed and efficiency when executing a level 3 dense linear algebra processing on a computer, said method comprising:

automatically setting an optimal machine state on said computer for said processing by selecting an optimal matrix subroutine from among a plurality of matrix subroutines stored in a memory that could alternatively perform a level 3 matrix multiplication

processing, wherein said computer includes an L1 cache, said method further comprising:

determining a size of each of matrices involved in said matrix multiplication;

and

selecting one of said matrices to reside in an L1 cache, based on said determined size,

wherein said selecting a matrix subroutine comprises determining which of said matrix subroutines is consistent with said matrix selected to reside in said L1 cache.

2. (Canceled)

3. (Previously presented) The method of claim 1, wherein said matrix subroutine comprises a substitute of a subroutine from LAPACK (Linear Algebra PACKage).

4. (Previously presented) The method of claim 3, wherein said substitute LAPACK subroutine comprises a Basic Linear Algebra Subroutine (BLAS) Level 3 L1 cache kernel.

5. (Currently amended) ~~The method of claim 1~~ A method of improving at least one of speed and efficiency when executing a level 3 dense linear algebra processing on a computer, said method comprising:

automatically setting an optimal machine state on said computer for said processing by selecting an optimal matrix subroutine from among a plurality of matrix subroutines stored in a memory that could alternatively perform a level 3 matrix multiplication processing, wherein said selecting a matrix subroutine comprises an aspect of a generalized matrix streaming process in which matrix data is stored in multiple levels of computer memory, including a matrix block stored in an L1 cache and matrix data of two other matrices stored in at least one higher level of cache, such that said matrix data of said two other matrices is systematically streamed into said matrix multiplication processing through said L1 cache.

6. (Previously presented) The method of claim 1, wherein said plurality of matrix subroutines comprises six possible matrix subroutines that could alternatively be used for said level 3 matrix multiplication processing.

7. (Currently amended) An apparatus, comprising:

a memory to store matrix data to be used for a processing in a level 3 dense linear algebra program;
an L1 cache;
a processor to perform said processing; and
a selector to select an optimal one of a plurality of possible matrix subroutines to that could alternatively perform said processing, thereby automatically setting said apparatus into

an optimal machine state to perform said processing, wherein said selector makes the selection by:

determining a size of each of matrices involved in said level 3 processing; and selecting one of said matrices to reside in said L1 cache, based on said determined sizes,

wherein said selecting a matrix subroutine comprises determining which of said matrix subroutines is consistent with said matrix selected to reside in said L1 cache.

8. (Canceled)

9. (Previously presented) The apparatus of claim 7, wherein said matrix subroutine comprises a substitute of a subroutine from LAPACK (Linear Algebra PACKage).

10. (Previously presented) The apparatus of claim 9, wherein said substitute LAPACK subroutine comprises a Basic Linear Algebra Subroutine (BLAS) Level 3 L1 cache kernel.

11. (Currently amended) The apparatus of claim 7 An apparatus, comprising:

a memory to store matrix data to be used for a processing in a level 3 dense linear algebra program;
a processor to perform said processing; and
a selector to select an optimal one of a plurality of possible matrix subroutines to that could alternatively perform said processing, thereby automatically setting said apparatus into an optimal machine state to perform said processing, wherein said selector for selecting a matrix subroutine includes a storage for storing matrix data in multiple levels of computer

memory and a mechanism for streaming said matrix data into said matrix multiplication process.

12. (Original) The apparatus of claim 7, wherein said plurality of matrix subroutines comprises six possible matrix subroutine kernel types.

13-20. (Canceled)